arXiv:1512.08103v1 [cs.CV] 26 Dec 2015

Data Driven Robust Image Guided Depth Map Restoration

Wei Liu[1], Yun Gu[1], Chunhua Shen[2], Xiaogang Chen[3], Qiang Wu[4] and Jie Yang[1]
[1]Shanghai Jiao Tong University, Shanghai, China. {liuwei.1989,geron762,jieyang}@sjtu.edu.cn
[2]University of Adelaide, Adelaide, Australia. chunhua.shen@adelaide.edu.au
[3]University of Shanghai for Science and Technology, Shanghai, China. xg.chen@live.com
[4]University of Technology, Sydney, Australia. Qiang.Wu@uts.edu.au

Abstract

Depth maps captured by modern depth cameras such as Kinect and Time-of-Flight (ToF) are usually contaminated by missing data, noises and suffer from being of low resolution. In this paper, we present a robust method for high-quality restoration of a degraded depth map with the guidance of the corresponding color image. We solve the problem in an energy optimization framework that consists of a novel robust data term and smoothness term. To accommodate not only the noise but also the inconsistency between depth discontinuities and the color edges, we model both the data term and smoothness term with a robust exponential error norm function. We propose to use Iteratively Reweighted Least Squares (IRLS) methods for efficiently solving the resulting highly non-convex optimization problem. More importantly, we further develop a data-driven adaptive parameter selection scheme to properly determine the parameter in the model. We show that the proposed approach can preserve fine details and sharp depth discontinuities even for a large upsampling factor ($8\times$ for example). Experimental results on both simulated and real datasets demonstrate that the proposed method outperforms recent state-of-the-art methods in coping with the heavy noise, preserving sharp depth discontinuities and suppressing the texture copy artifacts.

1. Introduction

Acquisition of depth information of 3D scenes is essential for many applications in computer vision and graphics. Applications range from 3D modeling to 3DTV and augmented reality. A number of applications require accurate and high-resolution depth maps, for instance, object reconstruction, robot navigation and automotive driver assistance. Recently, modern depth cameras such as Kinect and Time-of-Flight (ToF) (e.g, SwissRanger SR4000) shown impressive results and become increasingly affordable. They can obtain dense depth measurements at a high frame rate.

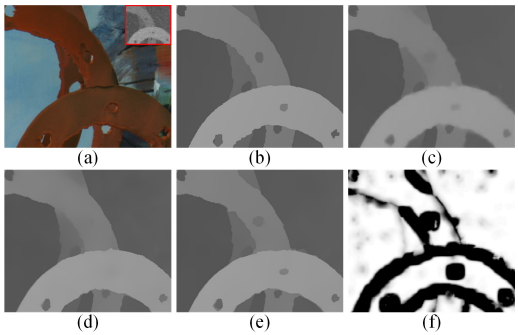

Figure 1. $8\times$ upsampling results. (a) The noisy low resolution depth map patch and the corresponding color image. (b) The ground truth. (c) The upsampling result of the state-of-the-art method of [31]. The upsampling result of our method (d) without adaptive bandwidth selection and (e) with adaptive bandwidth selection. (f) The corresponding bandwidth map of our adaptive bandwidth selection.

However, their depth maps usually suffer from missing values, noise and being of low resolution.

To facilitate the use of depth data, tremendous efforts have been spent on the restoration of depth maps obtained by modern depth cameras. The depth map can be restored by different example-based methods such as [22, 17, 10, 7], which enhance the quality with a single depth map. This category of methods tend to fail to cope with large upsampling factors and most state-of-the-art methods mainly focus on up to $4\times$ upsampling. Another direction is to restore the depth map from multiple low-quality depth maps such as [26, 8, 25, 11]. This category of methods are more practical for static scenes than for dynamic environments. There are also strong research interests in developing image guided restoration schemes such as [5, 14, 24, 31, 32, 18], which restore the depth map with the guidance of the registered (aligned) color image. These methods are often based on the assumption that there exists a joint occurrence between depth discontinuities and color image edges. They can produce promising restoration quality with larger upsampling factors and are also not subject to static scenes

when compared with the first two categories of methods. However, their depth maps often suffer from texture copy artifacts and blurring depth discontinuities when the depth discontinuities are inconsistent with the color edges.

In this paper, we propose a novel technique for robust image guide depth map restoration. The main contributions of our paper are as follows.

1. To handle heavy noises in depth maps, we develop a robust data term that measures 'pixel to patch' difference which is penalized by a robust error norm function within a Gaussian window. The proposed data term is inspired by the work in image denoising [23] and image editing [1]. We show that it is more robust in presence of heavy noises. To our knowledge, we are the first to introduce this robust data term in guided depth map restoration, and achieve impressive results.
2. The proposed color image guided smoothness term is robust against the inconsistency between the depth discontinuities and the color edges. It has been proven to perform well in suppressing the texture copy artifacts and offer much better performance in preserving sharp depth discontinuities than the state-of-the-art methods such as the auto-regressive model [31].
3. While the proposed optimization framework is highly non-convex, we propose a numeric solution with an Iteratively Re-weighted Least Squares (IRLS) optimization framework, which can be efficiently solved and easily implemented by solving a linear system using off-the-shelf linear system solvers.
4. More importantly, to eliminate heuristic parameter selection, we present a data-driven scheme to properly determine the parameter in our model such that fine details and sharp depth discontinuities are well preserved even for a large upsampling factor such as $8\times$, thus making the proposed method more practical.

Experimental results on simulated and real data show that our method outperforms state-of-the-art methods in terms of both visual quality and accuracy.

Related work and practical issues As our method belongs to the image guided restoration technique, we mainly review some recent guided restoration techniques and then the practical issues faced by this category of methods. Image guided depth restoration methods can be roughly classified as *local methods* and *global methods*. Local methods are based on a certain edge preserving filter where each pixel in the restored depth map is a weighted sum of its neighbors in the input depth map. The Joint Bilateral Upsampling (JBU) [14] extended the bilateral filter [29] for depth upsampling where the bilateral weight is computed on the guidance color image. JBU was also used to fill the holes for Kinect depth map restoration in [30, 31]. To achieve sharper depth discontinuities, the work in [18] evaluated pixel dissimilarities based on the geodesic distance instead of the Euclidean distance in the JBU. Very recently, a new edge preserving filter named Guide Filter (GF) [9] was introduced and used to perform joint upsampling. Their work was further improved in [20, 28] to preserve sharp depth discontinuities and spatial variation. Global methods usually formulate the restoration as optimization. The restored depth map is the minimal of the energy function. The authors of [5] performed restoration using an Markov Random Field (MRF) formulation with a pairwise appearance consistency data term and an image guided smoothness term. Their work was further extended in [24] by incorporating different weighting schemes to combine different cues including segmentation, image gradients, edge saliency and a non-local means. The work in [30, 31] performed image guided depth map restoration using an color guided auto-regressive model. Sparse representation model [12, 16] learns statistical dependencies between intensity and depth in a scene.

The fundamental assumption of image guided depth map restoration is that there exists a joint occurrence between depth discontinuities and color image edges. However, when the depth discontinuities are inconsistent with the color edges, these methods may face the following two problems: 1) texture copy artifacts on the smooth depth region when the corresponding color image is highly textured; 2) blurring depth discontinuity when the corresponding color image is more homogeneous. To handle these two issues, most recent work proposed to take the bicubic interpolation of the noisy low resolution depth map into account such as the definition of the auto-regressive coefficient in [31] and the RGB-D structure similarity in [16]. However, the bicubic interpolation of the input depth map becomes unreliable especially when the upsampling factor is large (e.g., $8\times$) and the input depth map contains heavy noises. This situation is often the case as the resolution of modern depth cameras (e.g., SwissRanger SR4000) is typically low and also contains heavy noises. We attempt to alleviate these two issues in this paper. The numeric solution of the proposed model shows that our model can iteratively exploit the newly updated depth map at each iteration. The quality of the newly updated depth map is considerably better than the bicubic interpolation of the input depth map. We show that this helps to preserve much sharper depth discontinuities than the previous work such as the auto-regressive model [31]. Also, how to well handle the heavy noises in the depth map is another important issue. Nevertheless, little work has been done on this issue. Here we demonstrate the robustness of our energy minimization model, due to the newly proposed robust data term.

2. The proposed method

The model Our upsampling model consists of two terms: the data term and the smoothness term. Given a noisy low

resolution depth map D_L, it is first interpolated to D^0 by bicubic interpolation. D^0 has the same resolution as the guidance image. Then our upsampling model is formulated as:

$$D_H = \arg\min_D \left\{(1-\alpha)E_D(D, D^0) + \alpha E_S(D)\right\} \quad (1)$$

where $E_D\left(D, D^0\right)$ is the data term that makes the result to be consistent with the input. $E_S\left(D\right)$ is the smoothness term that reflects prior knowledge of the smoothness of our solution. The relative importance of these two terms is balanced with the parameter α.

The data term $E_D\left(D, D^0\right)$: the data term is defined as:

$$E_D\left(D, D^0\right) = \sum_{i\in\Omega}\sum_{j\in N(i)} \omega_{i,j}\varphi_D(|D_i - D_j^0|^2) \quad (2)$$

where Ω represents the set of the high resolution coordinates. $N(i)$ is the neighborhood of i, namely, the square patch of radius r centered at i. The Gaussian window $\omega_{i,j}$ decreases the weights when j is far from i:

$$\omega_{i,j} = \exp\left(-\frac{|i-j|^2}{2\sigma_s^2}\right) \quad (3)$$

where σ_s is a constant that is defined by the user. $\varphi_D(\cdot)$ is the robust error norm function that we denote as the exponential error norm:

$$\varphi_D(x^2) = 2\lambda^2\left(1 - \exp\left(-\frac{x^2}{2\lambda^2}\right)\right) \quad (4)$$

where λ is a user defined constant.

The proposed data term is inspired by the work in image denoising [23] and image editing [1]. Their data term replaces the 'pixel to pixel' difference with 'pixel to patch' difference and has been shown robust against noisy input images (for image denoising [23]) and inaccurate input strokes (for image editing [1]). As depth maps obtained by a modern depth camera such as SwissRanger SR4000 are usually of low resolution and contaminated by heavy noises, we thus also employ the similar 'pixel to patch' difference measurement in the data term. The 'pixel to patch' difference may blur the depth discontinuities for the pixels around the depth discontinuities. To better preserve the depth discontinuities, we adopt the robust error norm function defined in Eq. (4) to model the data term instead of the L_2 error norm. This is because the error norm function used here is known to be robust against the outliers and thus could better preserve the depth discontinuities. The Gaussian window is introduced to further reduce the influence of the pixels far from the central pixel.

Our data term is different from that of previous methods [5, 24, 31] that only use 'pixel to pixel' difference modeled by the L_2 error norm. Their data term works well for input depth maps with high accuracy, but fails when heavy noises are presented in the data. In fact, our data term can be viewed as a generalized form of their data term. If we set the radius of the neighborhood $N(i)$ of pixel i as $r = 0$, then our data term becomes to the conventional non-robust form.

The smoothness term $E_S(D)$: Our smoothness is guided by the aligned color image. It is defined as:

$$E_S(D) = \sum_{i\in\Omega}\sum_{j\in N(i)} \omega_{i,j}^c\varphi_S\left(|D_i - D_j|^2\right). \quad (5)$$

The guide weight $\omega_{i,j}^c$ is defined as:

$$\omega_{i,j}^c = \omega_{i,j}\cdot\exp\left(-\frac{\sum_{k\in C}|I_i^k - I_j^k|^2}{3\times 2\sigma_c^2}\right) \quad (6)$$

where $\omega_{i,j}$ is the same as Eq. (3). $C = \{R, G, B\}$ represents the different channels of the color image. σ_c is a constant defined by the user. In fact, there are also other choice of the guide weight such as the shape-based structure aware weight in [31] or the more complex guide weight in [24]. However, we find that the adopted guide weight has already been enough to yield satisfying results. To adopt the weight in [31] and [24] shows little performance improvement while their computational cost is much higher than Eq. (6).

We also employ the function in Eq. (4) to model the smoothness term, i.e.:

$$\varphi_S(\cdot) = \varphi_D(\cdot). \quad (7)$$

As we will show in the sequel, the numeric solution to our model shows that the employed $\varphi_S(\cdot)$ makes the model robust against the inconsistency between the depth discontinuities and the color edges. This is significant in suppressing the texture copy artifacts and preserve sharp depth discontinuities.

2.1. The numeric solution and analysis

In this section, we first present the numeric solution to our formulated problem. Then we show further analysis why our model can handle the three practical issues discussed above, namely, the heavy noise in the input depth map, the texture copy artifacts and the blurring depth discontinuities caused by the inconsistency between the depth discontinuities and the color edges.

Due to the highly non-convex property of the proposed model, directly solving it is challenging. Previous work such as the one in [19] used the Loop Belief Propagation (LBP) [33] to solve their energy minimization function. However, classical energy minimization solvers such as the LBP [33] and graph cuts [2, 13] work for discrete energy minimization. In our problem, the variables are naturally real-valued and one has to discretize them in order

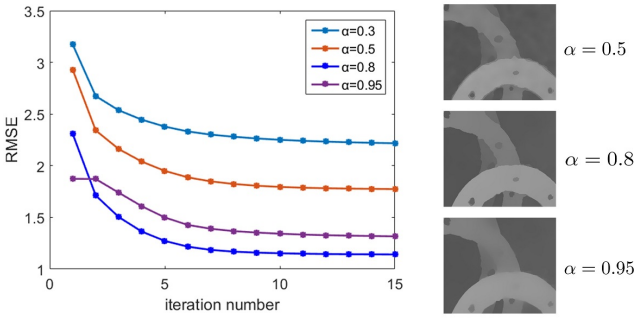

(a) (b)

Figure 2. (a) Convergency analysis of the proposed model on 8× upsampling for different α values in terms of RMSE between the newly updated depth map and the groundtruth in each iteration. (b) Example results of different α values.

to apply LBP or graph cuts. The computational cost can be extremely expensive when the continuous problem is discretized into 8192 levels (13 bits). This is the practical issue as modern depth cameras usually save the depth map in a 16 bits image (with 13 bits for depth values). The quantization error can be large if only a few quantization levels are used. In contrast, we present a numeric solution to our model that works for continuous system and can be efficiently solved. First, we present the normal equation of our model:

$$\begin{aligned}\frac{\partial E}{\partial D_i} &= (1-\alpha)\sum_{j\in N(i)}\omega_{i,j}d_{i,j}\left(D_i - D_j^0\right)\\ &+ 2\alpha\sum_{j\in N(i)}\omega^c_{i,j}s_{i,j}\left(D_i - D_j\right) = 0,\ i\in\Omega\end{aligned} \tag{8}$$

where we define

$$\begin{aligned}d_{i,j} = \varphi'_D(|D_i - D_j^0|^2),\quad s_{i,j} = \varphi'_S(|D_i - D_j|^2),\\ \varphi'_D(x^2) = \varphi'_S(x^2) = \exp\left(-\frac{x^2}{2\lambda^2}\right)\end{aligned} \tag{9}$$

$\varphi'_D\left(x^2\right) = \varphi'_S\left(x^2\right)$ is the derivative of $\varphi_D\left(x^2\right) = \varphi_S\left(x^2\right)$ defined in Eq. (4).

A closed-form solution to Eq. (8) is not available, we can only solve it iteratively. If we keep $s_{i,j}, d_{i,j}$ as constant in each iteration where $d^n_{i,j} = \varphi'_D(|D^n_i - D^0_j|^2), s^n_{i,j} = \varphi'_D(|D^n_i - D^n_j|^2)$ for iteration $n+1$, then Eq. (8) becomes the standard form of the following re-weighted least squares optimization framework as:

$$\begin{aligned}D^{n+1} = \arg\min_D\{(1-\alpha)\sum_{i\in\Omega}\sum_{j\in N(i)}\omega_{i,j}d^n_{i,j}(D_i - D^0_j)^2\\ +\alpha\sum_{i\in\Omega}\sum_{j\in N(i)}\omega^c_{i,j}s^n_{i,j}(D_i - D_j)^2\}\end{aligned} \tag{10}$$

Then we can iteratively solve Eq. (10) until the final output meets the convergence condition. This approximation is similar to the well-known Iteratively Re-weighted Least Squares (IRLS) [3] in the literature. However, their IRLS is only suitable for $L_p(0 < p < 2)$ norm optimization framework. In this paper, we also denote Eq. (10) as IRLS. As Eq. (10) is quadratic in each iteration, thus it can be minimized by solving the set of linear equations:

$$\begin{aligned}&\left[(1-\alpha)\sum_{j\in N(i)}\omega_{i,j}d^n_{i,j} + 2\alpha\sum_{j\in N(i)}\omega^c_{i,j}s^n_{i,j}\right]D_i\\ &- 2\alpha\sum_{j\in N(i)}\omega^c_{i,j}s^n_{i,j}D_j = (1-\alpha)\sum_{j\in N(i)}\omega_{i,j}d^n_{i,j}D^0_j\end{aligned} \tag{11}$$

we rewrite Eq. (11) in matrix notation as:

$$\begin{aligned}&[(1-\alpha)W^n - 2\alpha S^n]\,D = (1-\alpha)Z^nD^0\\ \Longrightarrow\ &D^{n+1} = (1-\alpha)\left[(1-\alpha)W^n - 2\alpha S^n\right]^{-1}Z^nD^0\end{aligned} \tag{12}$$

where D and D^0 are the vectors of the updated depth map and the initial depth map respectively. We use bicubic interpolation of the input depth map as initialization. W^n is a diagonal matrix with $W^n_{i,i} = \sum_{j\in N(i)}\omega_{i,j}d^n_{i,j} + \frac{2\alpha}{1-\alpha}\omega^c_{i,j}s^n_{i,j}$. S^n is the affinity matrix whose elements are $S^n_{i,j} = \omega^c_{i,j}s^n_{i,j}$. Z^n is another affinity matrix with $Z^n_{i,j} = \omega_{i,j}d^n_{i,j}$.

Solving the linear equation in Eq. (12) has long been studied in the literature and there are many efficient modern solvers. In our experiment, we use the Preconditioned Conjugate Gradient (PCG) in [15] to solve Eq. (12) which shows to be very efficient and can produce good results. Besides the efficiency, our IRLS also has good convergence property. Fig. 2(a) shows the convergence analysis for different α values. It is clear that our IRLS shows good convergence property for a large range of α values. Experiments on other upsampling factors and the Kinect dataset show similar performance. Note that the α value used in the our experiments is based on the consideration of noise smoothing and fine details preserving. Small α cannot well smooth the noise

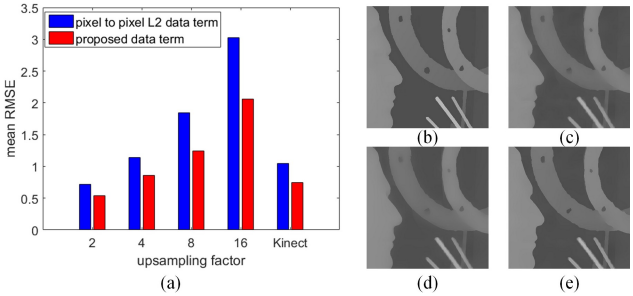

Figure 3. (a) Quantitative comparison of the proposed data term vs. the 'pixel to pixel' difference L_2 error norm data term in terms of RMSE on different upsampling factors of simulated ToF data [31] upsampling and simulated Kinect data [21] restoration. Examples of (b) the groundtruth, (c) results of pixel to pixel difference L_2 error norm data term, (d) first denoise the depth map with BM3D [4] and then upsample the depth map with the pixel to pixel difference L_2 error norm data term, (e) results of the proposed data term.

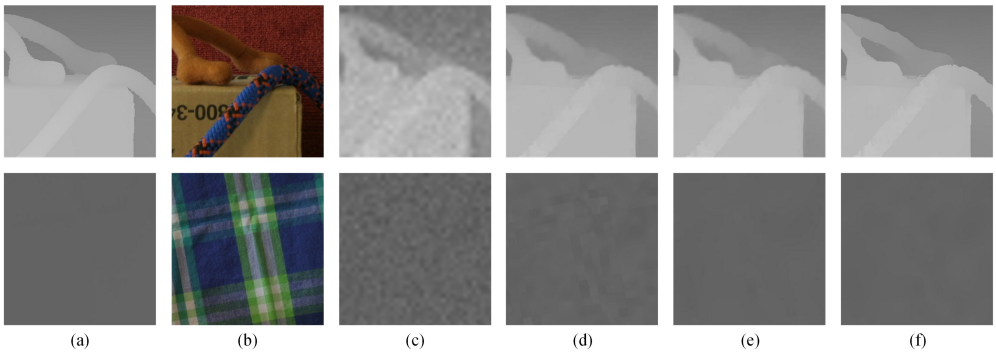

(a) (b) (c) (d) (e) (f)

Figure 4. Comparison of depth discontinuities preserving and texture copy artifacts suppression on $8\times$ upsampling. (a) The groundtruth. (b) The guidance color images. (c) The bicubic interpolation of the input depth map. Results obtained by the MRF [5], (e) the auto-regressive model [31] and (f) our method. The first row shows depth discontinuities preserving comparison. The second row shows texture copy artifacts suppression comparison.

while large α can over smooth fine details. Fig. 2(b) illustrates three example results of different α values.

Now we show how the proposed method may alleviate the three practical issues mentioned at the beginning of this section. First, note that on the right of Eq. (11) is $\sum_{j\in N(i)} \omega_{i,j} d^n_{i,j} D^0_j$ for each pixel $i \in \Omega$. This is in fact the filtered output of the initial depth map D^0 which results from our novel data term. If we use the data term that measures the L_2 norm of the pixel to pixel difference, then the right side of Eq. (11) will be only D^0_i for each pixel $i \in \Omega$. This will be quite noisy for noisy input and large upsampling factor. On the contrary, the filtering on the right side of Eq. (11) can well smooth the noise in D^0. This makes our model more robust against the noise. To further validate our analysis, we replace the proposed data term in our model with the pixel to pixel difference measurement L_2 error norm data term and perform experiments on both the simulated ToF dataset [31] and Kinect dataset [21]. Then we compare the mean RMSE of the results. Fig. 3 shows the comparison results. As shown in the figure, the proposed data term can clearly improve the performance especially for large upsampling factors. Fig. 3(c) and (e) also illustrate examples for visual comparison where the proposed data term also shows better visual quality.

In fact, the proposed data term is equivalent to the L_2 norm 'pixel to pixel' data term of which the input depth map is a filtered depth map as the right side of Eq. (11). However, this is different from the way that one firstly denoises the input depth map with a denoising method such as [4] and then performs upsampling with the L_2 norm 'pixel to pixel' data term. This is because our filtering weights $d^n_{i,j}$ is based on the newly updated depth map. Moreover, as pointed out in [22], first denoising the depth map followed by upsampling can destroy small structures in the depth map, especially for large upsampling factors. This is also validated in our experiments as illustrated in Fig. 3(d).

To handle the inconsistency between the depth discontinuities and the color edges, most recent work proposed to take the bicubic interpolation of the input depth map into account [31, 16]. This can efficiently suppress the texture copy artifacts but fail to properly preserve sharp depth discontinuities when the upsampling factor is large and the input depth map contains heavy noise. This is because the depth discontinuities have already been blurred for the bicubic interpolation of large upsampling factor as can be seen in Fig. 4(c). Note the guide weight of the smoothness term of our IRLS in Eq. (10) $s^n_{i,j}$ is based on the newly updated depth map of which the quality is much better than the bicubic interpolation. Fig. 4 shows the comparison of three different methods: 1) the Markov Random Field (MRF) in [5] which only use the color information for smoothness term weights, 2) the auto-regressive model in [31] which use both the color information and the bicubic interpolation of the input depth map, 3) our method which use both the color information and the newly updated depth map. It is clear that our method can not only suppress the texture copy artifacts but also preserve sharper depth discontinuities than the auto-regressive model in [31].

2.2. Data driven parameter selection

The parameter λ in Eq. (4) is an important parameter in our model. A large λ will have better noise smoothing but may over smooth the depth discontinuities. Thus, pixel in the homogeneous regions on the depth map should be assigned with large λ. A small λ could better preserve the depth discontinuities but performs poorly in noise smoothing. Thus, pixels along the depth discontinuities should be assigned with small λ. We denote λ as the bandwidth of the exponential error norm function. To eliminate heuristic parameter selection, we describe another optimization framework that adapts λ to each pixel in the depth map resulting in a data driven adaptive parameter selection model. Because the depth map is quite piece wise smooth, we assume

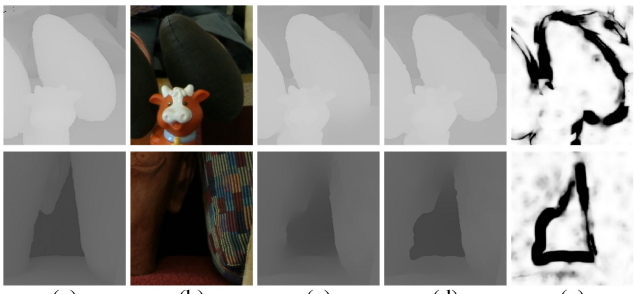

(a) (b) (c) (d) (e)

Figure 5. Visual comparison of our method for $8\times$ upsampling with and without bandwidth selection. (a) The groundtruth. (b) The corresponding color image. (c) Results obtained without bandwidth selection. (d) Results obtained with bandwidth selection and (e) the corresponding bandwidth maps.

that the bandwidth is also regular and smooth. Therefore, we add another term that consists of the L_2 norm of the gradient of $\lambda_i (i \in \Omega)$ to the objective function in Eq. (1), resulting in the following objective function:

$$E = (1-\alpha) E_D\left(D, D^0\right) + \alpha E_S(D) + \beta \sum_{i\in\Omega} |\nabla \lambda_i|^2,\ i \in \Omega \tag{13}$$

By minimizing Eq. (13) with respect to λ_i through the steepest gradient descent according to:

$$\lambda_i^{n+1} = \lambda_i^n - \tau \frac{\partial E}{\partial \lambda_i^n},\ i \in \Omega \tag{14}$$

where τ is the given updating rate and the derivative of the objective function is given by

$$\frac{\partial E}{\partial \lambda_i^n} =$$

$$(1-\alpha) \sum_{j\in N(i)} \omega_{i,j} \left[4\lambda_i^n \left(1 - d_{i,j}^n\right) - \frac{2\left(D_i^n - D_j^0\right)^2}{\lambda_i^n} d_{i,j}^n \right] +$$

$$\alpha \sum_{j\in N(i)} \tilde{\omega}_{i,j} \left[4\lambda_i^n \left(1 - s_{i,j}^n\right) - \frac{2\left(D_i^n - D_j^n\right)^2}{\lambda_i^n} s_{i,j}^n \right] + 2\beta \sum_{i\in\Omega} \Delta\lambda_i^n \tag{15}$$

where $d_{i,j}^n$ and $s_{i,j}^n$ are the same as Eq. (9).

In our experiments, depth map updating and the parameter selection are addressed in an iterative fashion through alternating the parameter update in Eq. (14) and the depth map update in Eq. (12). Fig. 1(f) illustrates a bandwidth map obtained by our method. It is clear that this bandwidth map well corresponds to the character of the depth map shown in Fig. 1(b). The adaptive selection around depth discontinuities corresponds to lower values than smooth homogeneous regions. Visual comparison illustrated in Fig. 1(d) and (e) shows that the proposed bandwidth selection helps to preserve fine details even for the $8\times$ upsampling. Fig. 5 further shows that the proposed bandwidth selection can also help to preserve sharp depth discontinuities even for $8\times$ upsampling.

3. Experiments

In this section, we compare our method with state-of-the-art methods for depth map restoration using the simulated ToF dataset from [31], simulated Kinect dataset from [21], the real ToF dataset from [6] and the NYU Kinect dataset [27]. All the experimental results in this paper and additional experimental results, including qualitative comparisons using other Kinect scene datasets, are provided in the supplemental material.

Experiments on simulated datasets We use the simulated ToF dataset from [31] and compare our method with the method in [24] which we denote as NLM-MRF, the image guided anisotropic total generalized variation upsampling in [6], the joint geodesic upsampling in [18], the joint intensity and depth co-sparse analysis model in [12] and the color guided auto-regression upsampling in [31][1]. The upsamling results are evaluated in root mean square error (RMSE) between the original depth map and the upsampling result. All the values of both the color image and the depth map are normalized into interval $[0, 1]$ for convenience. However, the RMSE is still reported in terms that all the values of the depth map are in interval $[0, 255]$. The parameters are set as follows: α in Eq. (1) is set as $0.7/0.75/0.8/0.9$ for $2\times/4\times/8\times/16\times$ upsampling. β in Eq. (13) is set to 0.5. The neighborhood $N(i)$ is chosen as a 9×9 square patch. σ_s and $\sigma_c$in Eq. (6) is set to 4 and $\frac{10}{255}$ respectively. The initial value of λ in Eq. (14) is set to $\frac{7}{255}$ for all $i \in \Omega$. Its updating rate τ in Eq. (14) is 0.3. We also reuse the simulated Kinect dataset from [21]. The parameter setting for this dataset is the same as that of the $8\times$ upsampling in the simulated ToF upsampling.

Table 1 summarizes the quantitative comparison in terms of RMSE. It is clear that our method outperforms all the compared methods especially for large upsampling factors $8\times$ and $16\times$. Fig. 6 further shows visual comparison with some $8\times$ upsampling examples. As highlighted, our method can well preserve sharp depth discontinuities and suppress the texture copy artifacts where depth discontinuities are inconsistent with color edges. Especially, our method can preserve sharper depth discontinuities than the auto-regressive model [31] which only use bicubic interpolation of the input depth map. Table 2 shows the mean RMSE comparison on the simulated Kinect dataset [21]. Our method also yields lowest mean RMSE on this dataset. Fig. 7 shows the visual comparison. Our method offers better performance in both smoothing the noise and preserving sharp depth discontinuities than other compared methods.

Experiments on real datasets To further test the pro-

[1]Their published MATLAB code requires huge amount of memory which is impractical for modern computers. The size of their results are also not the same as the guidance image. We thus implemented their code by ourselves with sparse matrixes and solve their normal equation with the same PCG [15] as ours.

Table 1. Quantitative comparison on the simulated ToF data using the noisy Middlebury dataset from [31]. Results are evaluated in RMSE and the best results are in bold.

	Art				*Book*				*Dolls*				*Laundry*				*Moebius*				*Reindeer*			
	2×	4×	8×	16×	2×	4×	8×	16×	2×	4×	8×	16×	2×	4×	8×	16×	2×	4×	8×	16×	2×	4×	8×	16×
JGF [18]	2.36	2.74	3.64	5.46	2.12	2.25	2.49	3.25	2.09	2.24	2.56	3.28	2.18	2.4	2.89	3.94	2.16	2.37	2.85	3.9	2.09	2.22	2.49	3.25
NLM-MRF [24]	1.66	2.47	3.44	5.55	1.19	1.47	2.06	3.1	1.19	1.56	2.15	3.04	1.34	1.73	2.41	3.85	1.2	1.5	2.13	2.95	1.26	1.65	2.46	3.66
JID [12]	1.69	2.98	3.68	5.99	1.53	2.71	3.04	4.39	1.54	2.71	2.94	3.90	1.45	2.72	3.16	4.63	1.55	2.72	2.94	4.34	1.65	2.80	3.13	4.63
TGV [6]	0.8	1.21	2.01	4.59	0.61	0.88	1.21	2.19	0.66	0.96	1.38	2.88	0.61	0.87	1.36	3.06	0.57	0.77	1.23	2.74	0.61	0.85	1.3	3.41
AR [31]	0.92	1.23	2.1	3.9	0.74	0.85	1.23	1.96	0.8	0.97	1.35	2.24	0.73	0.93	1.34	2.24	0.72	0.82	1.25	2.16	0.77	0.87	1.3	2.73
Ours	**0.68**	**1.07**	**1.69**	**3.01**	**0.54**	**0.8**	**1.1**	**1.65**	**0.61**	**0.85**	**1.15**	**1.78**	**0.6**	**0.83**	**1.17**	**2.01**	**0.53**	**0.76**	**1.11**	**1.74**	**0.55**	**0.82**	**1.12**	**2.13**

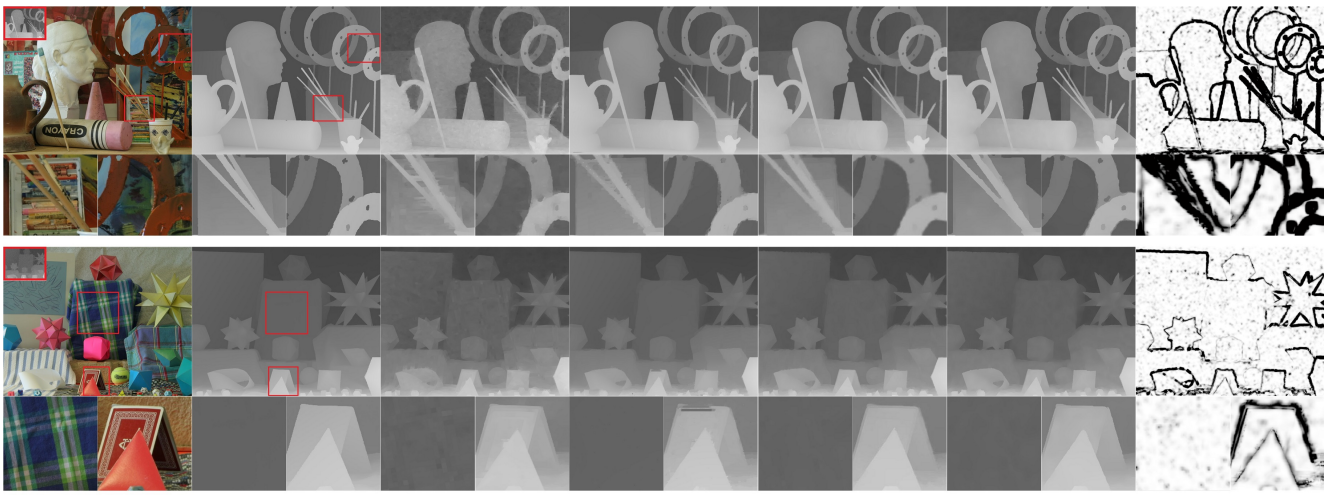

(a) (b) (c) (d) (e) (f) (g)

Figure 6. Visual comparison of 8× upsampling results on simulated ToF dataset from [31]. (a) The input depth maps (in red boxes) and the corresponding color images. (b) The groundtruth depth maps. The results of (c) the NLM-MRF in [24], (d) the image guided total generalized variation upsampling in [6], (e) the color guided auto-regressive model [31] and (f) our method and (g) the corresponding bandwidth maps by our bandwidth selection. Regions in red boxes are highlighted.

Table 2. Quantitative comparison on simulated Kinect data from [21]. Results are evaluated using the mean RMSE of all the data. The best results are in bold.

	NLM-MRF [24]	JID [12]	AR [31]	Ours
mean RMSE	1.27	2.15	1.02	**0.75**

Table 3. Quantitative comparison on real ToF dataset from [6]. The error is calculated as RMSE to the measured groundtruth in mm. The best results are in bold.

	JGF [18]	JBF [14]	NLM-MRF [24]	TGV [6]	AR [31]	Ours
Books	17.39mm	15.42mm	14.31mm	12.8mm	13.28mm	**12.33mm**
Devil	19.02mm	16.47mm	15.36mm	14.97mm	14.73mm	**14.12mm**
Shark	18.17mm	17.16mm	15.88mm	15.53mm	15.86mm	**14.71mm**

posed method, we also perform experiments on real datasets including the real ToF data set from [6] and the NYU Kinect dataset [27]. As summarized in table 3, our method also achieve lowest RMSE when compared with other state-of-the-art methods for ToF depth maps upsampling. Fig. 8 shows visual comparison with other methods. Our method performs better in suppressing the texture copy artifacts than others (enlarge the figure and see the 'ball' for clear comparison) and preserving sharp depth discontinuities. Fig. 9 shows the visual comparison on the Kinect data restoration. Our method has significant improvement over other methods on this dataset. Our method can properly smooth the noise, fill in the holes and rectify the jaggy depth discontinuities. Moreover, the depth discontinuities are quite sharp in our results. Results of the NLM-MRF [24] and color guided auto regressive model [31] suffer from blurring depth discontinuities. Results of the joint intensity and depth co-sparse analysis model [12] suffer from jaggy depth discontinuities.

4. Conclusion

In this paper, we propose a robust data driven optimization framework for high quality image guided depth map restoration. The newly proposed data term is robust against the noise which has been validated through both mathematical analysis and experimental results. The smoothness term performs well in suppress texture copy artifacts and performs especially better in preserving sharp depth discontinuities than other state-of-the-art methods. Given that our model is highly non-convex and directly solving it is challenging, we propose a numeric solution which can be efficiently solved and easily implemented by solving linear systems using off-the-shelf linear system solvers. The proposed data driven bandwidth selection further helps to preserve sharper depth discontinuities and fine details even

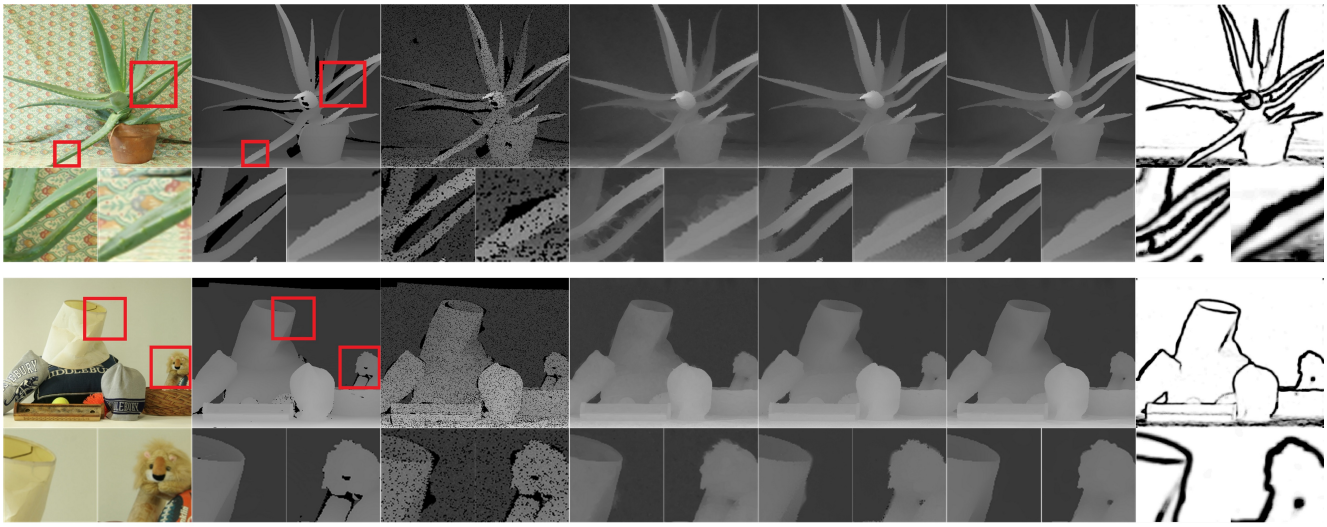

(a) (b) (c) (d) (e) (f) (g)

Figure 7. Visual comparison on simulated Kinect dataset from [21]. (a) The color images. (b) The groundtruth. (c) The degraded depth maps. The results of (d) the NLM-MRF in [24], (e) the color guided auto-regressive model [31], (f) our method and (g) the bandwidth maps by our bandwidth selection. Regions in red boxes are highlighted.

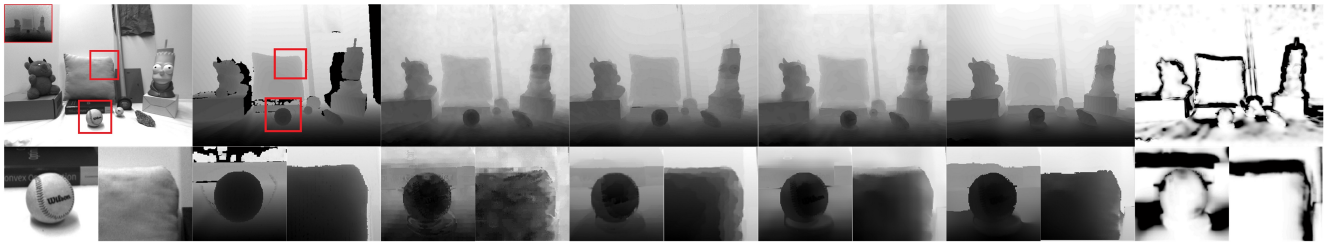

(a) (b) (c) (d) (e) (f) (g)

Figure 8. Visual comparison on real ToF dataset form [6]. (a) The input depth map (in the red box) and the corresponding color image. (b) The groundtruth depth map. The result of (c) the NLM-MRF in [24], (d) the image guided total generalized variation upsampling in [6], (e) the color guided auto-regressive model [31] and (f) our method and (g) the corresponding bandwidth map by our bandwidth selection. Regions in red boxes are highlighted.

for large upsampling factors. Experimental results on both simulated data and real data have shown our method outperforms other state-of-the-art methods.

References

[1] X. An and F. Pellacini. Appprop: all-pairs appearance-space edit propagation. *ACM Transactions on Graphics (TOG)*, 27(3):40, 2008. 2, 3

[2] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 23(11):1222–1239, 2001. 3

[3] R. Chartrand and W. Yin. Iteratively reweighted algorithms for compressive sensing. In *ICASSP*, pages 3869–3872. IEEE, 2008. 4

[4] K. Dabov, A. Foi, V. Katkovnik, and K. Egiazarian. Image denoising by sparse 3-d transform-domain collaborative filtering. *Image Processing, IEEE Transactions on*, 16(8):2080–2095, 2007. 4, 5

[5] J. Diebel and S. Thrun. An application of markov random fields to range sensing. In *NIPS*, 2005. 1, 2, 3, 5

[6] D. Ferstl, C. Reinbacher, R. Ranftl, M. Rüther, and H. Bischof. Image guided depth upsampling using anisotropic total generalized variation. In *ICCV*, 2013. 6, 7, 8

[7] D. Ferstl, M. Rüther, and H. Bischof. Variational depth superresolution using example-based edge representations. In *ICCV*, 2015. 1

[8] U. Hahne and M. Alexa. Exposure fusion for time-of-flight imaging. In *Computer Graphics Forum*, volume 30, pages 1887–1894. Wiley Online Library, 2011. 1

[9] K. He, J. Sun, and X. Tang. Guided image filtering. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(6):1397–1409, 2013. 2

[10] M. Hornacek, C. Rhemann, M. Gelautz, and C. Rother. Depth super resolution by rigid body self-similarity in 3d. In *CVPR*, 2013. 1

[11] S. Izadi and D. e. Kim. Kinectfusion: real-time 3d reconstruction and interaction using a moving depth camera. In

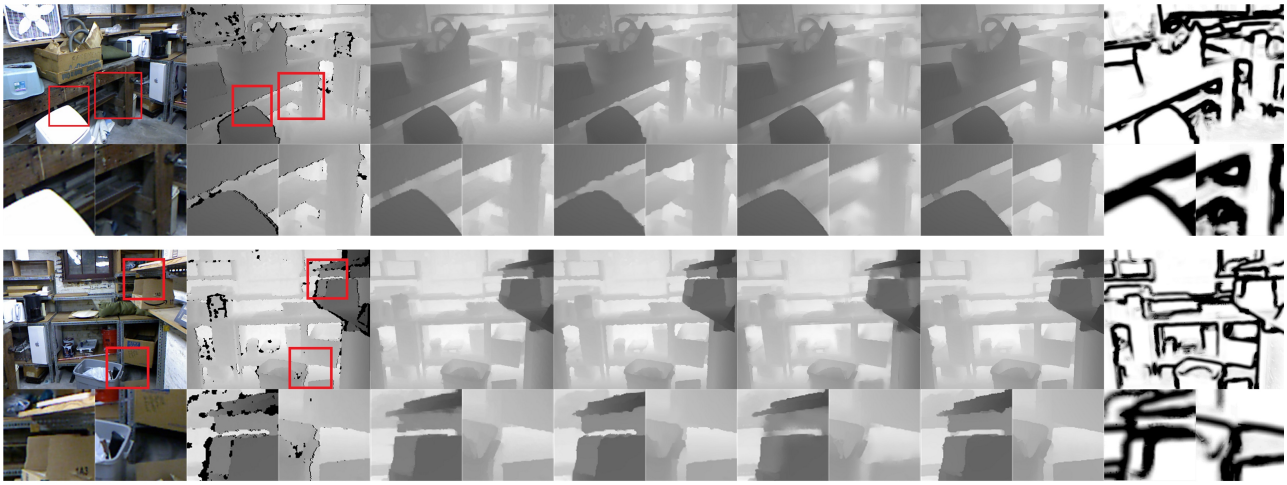

(a) (b) (c) (d) (e) (f) (g)

Figure 9. Visual comparison on NYU Kinect dataset from [27]. (a) The color images. (b) The RAW depth maps. The results of (c) the NLM-MRF in [24], (d) joint intensity and depth co-sparse analysis model in [12], (e) the color guided auto-regressive model [31], (f) our method and (g) the corresponding bandwidth maps by our bandwidth selection.

Proceedings of the 24th annual ACM symposium on User interface software and technology, pages 559–568. ACM, 2011. 1

[12] M. Kiechle, S. Hawe, and M. Kleinsteuber. A joint intensity and depth co-sparse analysis model for depth map super-resolution. In *ICCV*, pages 1545–1552. IEEE, 2013. 2, 6, 7, 9

[13] V. Kolmogorov and R. Zabin. What energy functions can be minimized via graph cuts? *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 26(2):147–159, 2004. 3

[14] J. Kopf, M. F. Cohen, D. Lischinski, and M. Uyttendaele. Joint bilateral upsampling. In *ACM Transactions on Graphics*, volume 26, page 96. ACM, 2007. 1, 2, 7

[15] D. Krishnan, R. Fattal, and R. Szeliski. Efficient preconditioning of laplacian matrices for computer graphics. *ACM Transactions on Graphics (TOG)*, 32(4):142, 2013. 4, 6

[16] H. Kwon, Y.-W. Tai, and S. Lin. Data-driven depth map refinement via multi-scale sparse representation. In *CVPR*, pages 159–167, 2015. 2, 5

[17] J. Li, Z. Lu, G. Zeng, R. Gan, and H. Zha. Similarity-aware patchwork assembly for depth image super-resolution. In *CVPR*, 2014. 1

[18] M.-Y. Liu, O. Tuzel, and Y. Taguchi. Joint geodesic upsampling of depth images. In *CVPR*, 2013. 1, 2, 6, 7

[19] J. Lu, D. Min, R. S. Pahwa, and M. N. Do. A revisit to mrf-based depth map super-resolution and enhancement. In *ICASSP*, pages 985–988. IEEE, 2011. 3

[20] J. Lu, K. Shi, D. Min, L. Lin, and M. N. Do. Cross-based local multipoint filtering. In *CVPR*, 2012. 2

[21] S. Lu, X. Ren, and F. Liu. Depth enhancement via low-rank matrix completion. In *CVPR*, pages 3390–3397. IEEE, 2014. 4, 5, 6, 7, 8

[22] O. Mac Aodha, N. D. Campbell, A. Nair, and G. J. Brostow. Patch based synthesis for single depth image super-resolution. In *ECCV*. 2012. 1, 5

[23] P. Mrázek, J. Weickert, and A. Bruhn. On robust estimation and smoothing with spatial and tonal kernels. In *Geometric properties for incomplete data*, pages 335–352. Springer, 2006. 2, 3

[24] J. Park, H. Kim, Y.-W. Tai, M. S. Brown, and I. Kweon. High quality depth map upsampling for 3d-tof cameras. In *ICCV*, 2011. 1, 2, 3, 6, 7, 8, 9

[25] A. Rajagopalan, A. Bhavsar, F. Wallhoff, and G. Rigoll. Resolution enhancement of pmd range maps. In *Pattern Recognition*, pages 304–313. Springer, 2008. 1

[26] S. Schuon, C. Theobalt, J. Davis, and S. Thrun. Lidarboost: Depth superresolution for tof 3d shape scanning. In *CVPR*, 2009. 1

[27] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus. Indoor segmentation and support inference from rgbd images. In *ECCV*, pages 746–760. Springer, 2012. 6, 7, 9

[28] X. Tan, C. Sun, and T. D. Pham. Multipoint filtering with local polynomial approximation and range guidance. In *CVPR*, pages 2941–2948. IEEE, 2014. 2

[29] C. Tomasi and R. Manduchi. Bilateral filtering for gray and color images. In *ICCV*, 1998. 2

[30] J. Yang, X. Ye, K. Li, and C. Hou. Depth recovery using an adaptive color-guided auto-regressive model. In *ECCV*. 2012. 2

[31] J. Yang, X. Ye, K. Li, C. Hou, and Y. Wang. Color-guided depth recovery from rgb-d data using an adaptive autoregressive model. *IEEE Transactions on Image Processing*, 23(8):3443–3458, 2014. 1, 2, 3, 4, 5, 6, 7, 8, 9

[32] Q. Yang, R. Yang, J. Davis, and D. Nistér. Spatial-depth super resolution for range images. In *CVPR*, 2007. 1

[33] J. S. Yedidia, W. T. Freeman, Y. Weiss, et al. Generalized belief propagation. In *NIPS*, volume 13, pages 689–695, 2000. 3